UNIVERSITÄT PASSAU
Fakultät für Mathematik und Informatik

Basic-block Graphs: Living Dinosaurs?

Jens Knoop[1], Dirk Koschützki, Bernhard Steffen
Fakultät für Mathematik und Informatik
Universität Passau
Innstraße 33
D–94032 Passau

{knoop,koschuet,steffen}@fmi.uni-passau.de

MIP-9715
Juli 1997

[1]The work of the author was partially funded by the Leibniz Programme of the German Research Council (DFG) under grant Ol 98/1-1.

Abstract

Since decades, *basic-block* (*BB*) graphs are the state-of-the-art means for representing programs in advanced industrial compiler environments. The usual justification for introducing the intermediate BB-structures in the program representation is *performance*: analyses on BB-graphs are generally assumed to outperform their counterparts on *single-instruction* (*SI*) graphs, which, undoubtedly, are conceptually much simpler, easier to implement, and more straightforward to verify.

In this article, we discuss the difference between the two program representations and show by means of runtime measurements that, according to the new computer generations, performance is no longer on the side of the more complex BB-graphs. In fact, it turns out that no sensible reason for the BB-structure remains. Rather, we will demonstrate that *edge-labeled* SI-graphs, which in contrast to the classical flow graphs model statements in their edges instead of in their nodes, are most adequate, both for the theoretical reasoning about and for the implementation of analysis and optimization algorithms. We are convinced that this perception has far-reaching consequences for the design of compiler systems.

Keywords

Basic blocks, basic-block graphs, data-flow analysis, efficiency, flow graphs, node-labeled vs. edge-labeled flow graphs, program optimization, single-instruction graphs.

Contents

List of Figures

List of Tables

1 Motivation

In program analysis and optimization it is common to work on so-called *flow graphs*, whose edge structure makes the control flow of the underlying program explicit. Most widely used are *node-labeled basic-block* (*BB*) graphs, whose nodes represent maximal sequences of straight-line code. This most prominent representation can be modified according to (1) its granularity in order to arrive at *single-instruction* (*SI*) graphs and/or to (2) its way of instruction modelling: *edge-labeled* graphs model instructions or basic blocks by edges rather than nodes.

In this article we investigate these four variants of program representation both from a theoretical and practical point of view. It turns out that the most prominent representation in practice is no longer adequate in times where already the main storage of home computers easily accommodates SI-graphs for huge procedures. Moreover, we show that *edge-labeled* graphs simplify both the theoretical reasoning about analysis and optimization as well as their implementation in comparison to their *node-labeled* counterparts. This is mainly due to the fact that edge-labeled graphs allow us to use nodes as the natural place for storing analysis results and information within the graph structure, whereas node-labeled graphs require separate means and operational overhead. The advantage of edge-labeled graphs is even more drastic when looking at programs with a parallel operator [22].

Our investigation is complemented by runtime measurements demonstrating that the "classical" reason for BB-structuring, i.e., opening analysis and optimization to realistic programs, did not survive the radical changes at the hardware front: as SI-graphs fit into main memory now, BB-graphs can at the most gain some performance for programs with basic blocks of comparatively large average size. In fact, in everyday's life, we never experienced any situation, where BB-graphs were superior to SI-graphs.

Moreover, the BB-structuring is limited in its application scenario (cf. Section 2.2). Thus, there are strong reasons to considering edge-labeled SI-graphs as the most adequate uniform representation for compiler optimization. In particular, it is fair to state that BB-graphs outlived their time, and that they can be considered *living dinosaurs*.

Structure of the Article: In Section 2 we critically re-investigate the properties usually attributed to BB-graphs. This leads directly to the central thesis of this article stating the superiority of edge-labeled SI-graphs for analysis and optimization. Subsequently, we present our preliminaries including a taxonomy of flow-graph variants in Section 3. Central are then Sections 4 and 5. In Section 4 we give theoretical evidence for the superiority of edge-labeled SI-graphs taking three different view-points. In Section 5 we complement this by runtime measurements demonstrating that BB-graphs do not compensate for their conceptual complexity in practice. Together, this confirms our thesis of Section 2 both theoretically and practically. Section 6, finally, contains our conclusions.

2 Basic Blocks: "Folk Knowledge"

2.1 Benefits

The central benefits commonly claimed are summarized by two keywords:

- *Performance:* ... because less nodes take part in costly fixed-point iterations.
- *Compactness:* ... because larger programs fit into the main memory.

Both points do not reflect the situation of the late nineties: state-of-the-art fixpoint algorithms can easily deal with graphs of more than 10^5 nodes in real time, a size which will hardly be exceeded by procedural SI-graphs (of course, this also depends on the fact that these graphs fully fit into the main memory of modern computers).

2.2 Short-comings

In contrast, BB-graphs are infected with a number of inquestionable short-comings:

- *Higher conceptual complexity:* ... basic blocks introduce undesired *hierarchy* complicating both theoretical reasonings as well as implementations (cf. Section 4.1)!
- *Demand for pre- and postprocesses:* ... usually required for managing the subtleties of hierarchy (e.g, *dead code elimination*, *constant propagation*, ...), or "tricky" formulations mandatory for by-passing them (e.g., *partial redundancy elimination*) (cf. Section 4.2)!
- *Limited generality:* ... certain practically relevant analyses and optimizations are hard or even impossible to be expressed on the basic-block level (e.g., *faint variable analysis and elimination*) (cf. Section 4.3)!

Higher conceptual complexity: Basic blocks structure a graph *hierarchically.* As a consequence, analysis and optimization problems must be designed, reasoned about, and implemented on two different levels, the *basic-block* and the *instruction* level; the latter in order to push the data-flow information computed globally for basic blocks to their constituting instructions. This two-level approach is particularly cumbersome, whenever the local analyses for several global analyses are performed in a single traversal, a situation which is common in practice for performance reasons. Maintaining a consistent view on basic blocks becomes then often a nontrivial task due to intricate interdependencies of different analyses and transformations based thereof (cf. [32]). This is a pity, particularly, because a single level, even more, the intellectually less sophisticated and less challenging instruction level would suffice. An observation, which previously was made by other researchers as well (cf. [32]). Here, however, we investigate its consequences in more detail and complement them with runtime measurements showing that the higher conceptual complexity of BB-graphs does not pay-off in practice in terms of performance.

Demand for pre- and postprocesses: Working on BB-graphs requires usually pre- and postprocesses on the analysis and optimization side. In fact, this holds for almost every optimizing program transformation, and is another source of additional conceptual complexity. Obvious, though comparatively simple examples are procedures for *dead code elimination* and *constant propagation*. After computing the required data-flow information for basic blocks, they must be inspected themselves by a postprocess in order to apply the transformation under consideration to the complete program. Sometimes pre- and postprocesses can be avoided. However, this usually relies on "tricky" formulations often injuring conceptual clarity and transparency of the transformation. A representative example is the *busy-code-motion* (*BCM*)-transformation of [20] for the elimination of partially redundant computations in a program. The *BCM*-transformation does not require a postprocess as e.g. dead code elimination. This, however, comes at the price of a more complicated reasoning about the correctness of the transformation as the *meet-over-all-paths* (*MOP*)-solutions of the data-flow properties involved do not directly fit to the *maximal-fixed-point* (*MFP*)-solutions computed because they apply to basic-block internal program points rather than to their "natural" entry and exit points (see Figure 4; note further the difference between the equation systems for *availability* (Section 4.1) and *very busyness* (Appendix A), and their counterparts for *up-safety* and *down-safety* involved in the *BCM*-transformation (Section 4.2)).

Limited generality: The applicability of BB-graphs for practically relevant problems is limited. *Faint code elimination* (cf. [11, 13, 21]), a generalization of dead code elimination, is a typical representative of such a problem, which seems to be impossible to be formulated on a basic-block level. The point is that the local properties of a basic block are not invariant under the global faintness analysis. This invariance, however, is the prerequisite for lifting an analysis from the instruction to the basic-block level, i.e., for hierarchically decomposing it into a global analysis on basic blocks followed by their local inspection.

2.3 What is Left?

Whereas modern computers easily deal with the size of even large SI-graphs, humans will hardly be able to comprehend small ones of a few hundred nodes only. Here a factor of 5 to 10 in size may well make a significant difference: it is easy to graphically deal with up to 60 or 80 nodes, but 500 nodes are definitely beyond a comfortable graphical treatment. Thus, the basic-block structure can well be regarded as a means to extend the range of graphically manageable programs.

Even though BB-graphs are a means to support graphical management, the question remains whether they are most adequate. Here the answer is no for two reasons:

- *Syntactic* reduction in terms of a macro or sub-routine concept, structuring the argument program, is much superior to an algorithmic BB-reduction, as this structuring allows an almost arbitrary reduction, while at the same time expressing some of the intention of the programmer. Thus the reduced programs "are meant" to be understandable.

- *Semantic* reduction according to a certain aspect of the program reduces the program by hiding all details irrelevant for the aspect currently under investigation. This reduction typically has an effect far beyond a BB-collapse, and it collapses program parts according to their properties rather than according to some comparatively trivial syntactic criterion. This allows to maintain understandability on the level of the collapsed program, which is by no means guaranteed by a BB-collapse.

Both syntactic and semantic reduction can easily be computed in real time in order to provide the user with the most adequate "view" of the program. Both reduction techniques have in fact been successfully applied in an industrial project, where they were one of the key "unique selling propositions" [30, 31].

In the remainder of this article we give theoretical and empirical evidence advocating our thesis that *edge-labeled SI*-graphs are the graph variant simultaneouly fitting the needs of theoreticians and practitioners best!

3 Preliminaries: A Taxonomy of Flow Graphs

Programs are basically represented by *directed flow graphs* consisting of a set of *nodes* and *edges* together with a unique *start node* and *end node*, which are assumed to have no incoming and outgoing edges, respectively. Flow graphs can either be *node-labeled* or *edge-labeled*; they can be *basic-block* (*BB*) graphs or *single-instruction* (*SI*) graphs. Together this leads to the taxonomy of flow graphs displayed in Figure 1. We recall that node-labeled BB-graphs are prevailing both in practice and in the literature on analysis and optimization. They can be considered the de-facto standard.[2] In contrast, we argue that from both a theoretical and practical point of view edge-labeled SI-graphs are the most appropriate flow-graph variant. Figure 2 illustrates the different flow-graph variants by means of small flow-graph fragments.

Conventions: We denote BB-graphs and SI-graphs by quadruples $\mathbf{G} = (\mathbf{N}, \mathbf{E}, \mathbf{s}, \mathbf{e})$ and $G = (N, E, s, e)$, respectively. Basic blocks are usually denoted by β, and instructions by ι, both possibly indexed. $lhs(\iota)$ denotes the left-hand side variable of an instruction ι, and $block(\iota)$ denotes the basic block containing ι. Moreover, $start(\beta)$ and $end(\beta)$ denote the first and the last instruction of β, respectively. For an SI-graph G, $pred_G(n) =_{df} \{ m \mid (m, n) \in E \}$ and $succ_G(n) =_{df} \{ m \mid (n, m) \in E \}$ denote the set of all immediate predecessors and successors of a node n. Sometimes, we will use ε as an identifier for edges. A *finite path* in G is a sequence $(n_1, \ldots, n_q)$ of nodes such that $(n_j, n_{j+1}) \in E$ for $j \in \{1, \ldots, q-1\}$. $\mathbf{P}_G[m, n]$ denotes the set of all finite paths from m to n, and $\mathbf{P}_G[m, n[$ the set of all finite paths from m to

[2] See e.g. [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 12, 13, 14, 16, 17, 20, 21, 23, 24, 25, 26, 27, 28, 32, 33, 34]. One of the few exceptions is [22] considering edge-labeled SI-graphs.

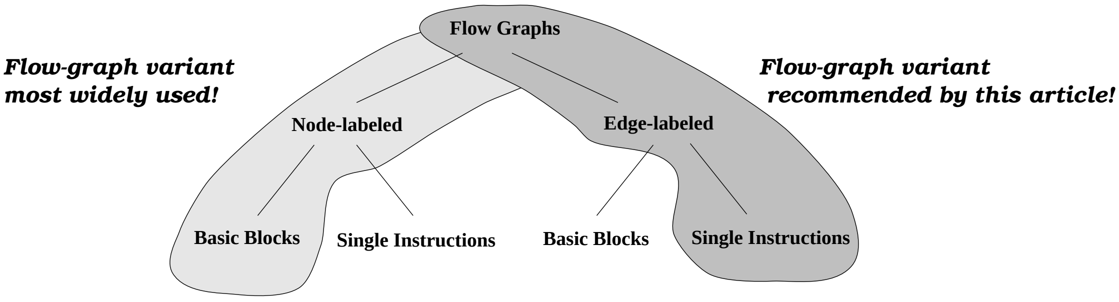

Fig. 1: A taxonomy of flow graphs.

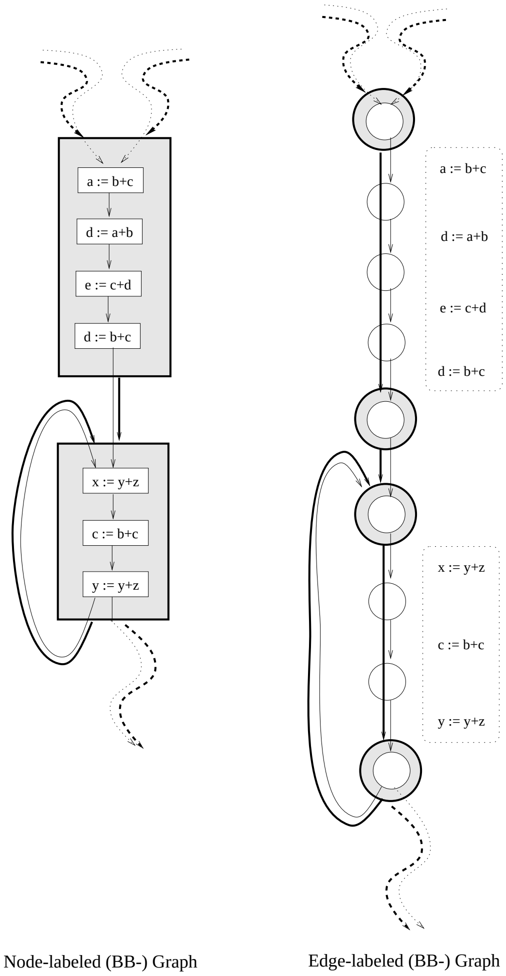

Fig. 2: Node-labeled vs. edge-labeled (BB- and SI-) flow graphs.

a predecessor of n. Finally, every node $n \in N$ is assumed to lie on a path from s to e. These notions are used analogously for BB-graphs.

4 Theory: Short-comings of BB-Graphs

In this section we give theoretical evidence for the short-comings of BB-graphs for analysis and optimization as summarized in Section 2.2. Each of the points mentioned, *higher computational complexity*, *demand for pre- and postprocesses*, and *limited generality*, is investigated in a separate subsection.

4.1 Higher Conceptual Complexity

BB-graphs are inherently hierarchical. They have a two-level structure. This enlarges the conceptual and technical complexity of specifying analysis problems, and of reasoning about them as well as their implementations. We demonstrate this on two different levels of abstraction. First, on the level of the abstract-interpretation framework underlying data-flow analysis (DFA) (cf. Section 4.1.1). Second, on the *concrete* level of a typical and practically relevant DFA-problem, the availability of program terms (cf. Section 4.1.2).

4.1.1 Correctness and Precision: *MOP*-Solution and *MFP*-Solution

Fundamental for reasoning about correctness and precision of a DFA are the *meet-over-all-paths* (*MOP*) solution and the *maximal-fixed-point* (*MFP*) solution in the sense of Kam and Ullman [15]. Intuitively, the *MOP*-solution defines the desired solution of a DFA-problem. It directly reflects the operational semantics of a program because it is the "meet (intersection)" of all (data-flow) informations contributed by some program path reaching a specific program point. Unfortunately, this solution does in general not induce an effective computation procedure. Fortunately, this usually holds for the *MFP*-solution of a DFA-problem. As the counterpart to the *MOP*-solution, the *MFP*-solution is defined as the greatest solution of an equation system expressing consistency of an annotation of the program with (data-flow) informations, which, under certain side-constraints, can effectively be computed by an iterative fixpoint procedure. For each analysis, however, the *MFP*-solution must be proved precise or at least correct with respect to the *MOP*-solution. Though this reduces in practice to checking the premises of the well-known Coincidence and Safety Theorems of Kildall [18], and Kam and Ullman [15], the technical complexity of the definitions of the *MOP*- and *MFP*-solution, and hence of applying these theorems, varies significantly for the flow-graph variants of Figure 1. We demonstrate this by contrasting the definitions of the *MOP*- and *MFP*-solution for edge-labeled SI-graphs, leading to the most elegant and concise versions, with their counterparts for node-labeled BB-graphs, leading to the most complex ones. It is common to all variants that the specification of a DFA-problem consists basically of a local semantic functional $[\![\ ]\!]$ describing the effect of the instructions (basic blocks) in terms of a function on a complete lattice $\mathcal{C}$ representing the data-flow informations of interest, and a start information $c_{\mathbf{s}} \in \mathcal{C}$, which is assumed to be valid on calling the program under consideration. The effect of a program path is then defined as the effect of the sequential composition of its elements (i.e., nodes or edges).

A) *MOP*- and *MFP*-Solution for Edge-labeled SI-Graphs

Considering an edge-labeled SI-graph G, the local semantic functional $[\![\ ]\!]_\iota : E \rightarrow (\mathcal{C} \rightarrow \mathcal{C})$ specifying a DFA-problem defines for every edge e of G a function on $\mathcal{C}$. The index ι indicates that the local semantic functions define the effect of instructions, not of basic blocks. The solution of the *MOP*-approach is then given by:

The *MOP*-Solution: $\forall\, c_s \in \mathcal{C}\ \forall\, n \in N.\ MOP_{([\![\ ]\!]_\iota, c_s)}(n) =_{df} \sqcap \{\, [\![\, p\,]\!]_\iota(c_s) \mid p \in \mathbf{P}_G[s, n] \,\}$

Note that this definition is the formal counterpart to the informal definition given above. Next, we define the corresponding *MFP*-solution.

The *MFP*-Solution: $\forall\, c_s \in \mathcal{C}\ \forall\, n \in N.\ MFP_{([\![\ ]\!]_\iota, c_s)}(n) =_{df} \texttt{info}_{c_s}(n)$

where info_{c_s} denotes the greatest solution of the following equation system:

$$\mathbf{info}(n) = \begin{cases} c_s & \text{if } n = s \\ \sqcap \{\, [\![\, (m, n)\,]\!]_\iota(\mathbf{info}(m)) \mid m \in pred_G(n) \,\} & \text{otherwise} \end{cases}$$

The Coincidence Theorem 4.1 is the handle for proving the *MFP*-solution of a DFA-problem precise with respect to its *MOP*-solution (cf. [15, 18]).[3]

Theorem 4.1 (Coincidence Theorem)
The MFP-solution and the MOP-solution coincide, i.e.,

$$\forall\, c_s \in \mathcal{C}\ \forall\, n \in N.\ MFP_{([\![\ ]\!]_\iota, c_s)}(n) = MOP_{([\![\ ]\!]_\iota, c_s)}(n)$$

if the local semantic functions $[\![\, e\,]\!]_\iota$, $e \in E$, *are all distributive.*

B) *MOP*- and *MFP*-Solution for Node-labeled BB-Graphs

For comparison we now recall the definitions of the *MOP*- and *MFP*-solution for node-labeled BB-graphs. This requires a two-level approach. First, defining them on the basic-block level. Second, defining them on the instruction level. The two levels are illustrated in Figure 3, which in addition makes the usually implicit distinction between eNtry- and eXit-points for both instructions (N-I,X-I) and basic blocks (N-BB,X-BB) of a node-labeled graph explicit. Note that making the eNtry- and eXit-points explicit, the node-labeled BB-graph is simply a (complicated) coding of an edge-labeled (SI-) graph.

Level 1 – Basic-block Level: On this level we need a local semantic functional $[\![\ ]\!]_\beta : \mathbf{N} \rightarrow (\mathcal{C} \rightarrow \mathcal{C})$, which defines the effect of complete basic blocks, not just of single instructions (cf. Figure 3). In practice, this requires a preanalysis of every basic block being accomplished by some preprocess. The definition of the *MOP*-solution is then as follows. Note that it defines for every node an eNtry- and an

[3] If the local semantic functions are monotonic, the *MFP*-solution is a safe (correct) approximation of the *MOP*-solution, i.e.: $MFP_{([\![\ ]\!]_\iota, c_s)} \sqsubseteq MOP_{([\![\ ]\!]_\iota, c_s)}$ (Safety Theorem, [15]).

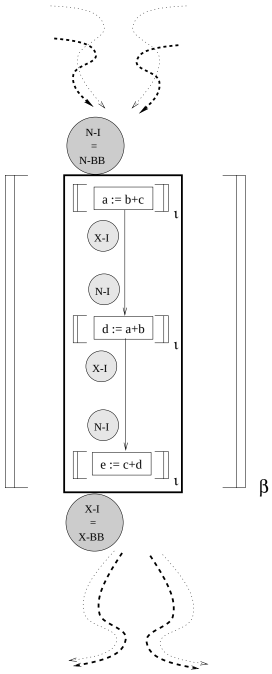

Fig. 3: Illustrating eNtry- and eXit-points and local semantic functions.

eXit-information, which is mandatory for node-labeled flow graphs. Unfortunately, this introduces an inhomogeneous notion of program point into the reasoning.

The *MOP*-Solution: (Basic-block Level)

$$\forall\, c_{\mathbf{s}} \in \mathcal{C}\ \forall\, \mathbf{n} \in \mathbf{N}.\ MOP_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathbf{n}) =_{df} (\, N\text{-}MOP_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathbf{n}),\ X\text{-}MOP_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathbf{n})\,)$$

with

$$N\text{-}MOP_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathbf{n}) =_{df} \bigsqcap \{\, [\![\, p\,]\!]_\beta(c_{\mathbf{s}}) \mid p \in \mathbf{P_G}[\mathbf{s}, \mathbf{n}[\,\}$$
$$X\text{-}MOP_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathbf{n}) =_{df} \bigsqcap \{\, [\![\, p\,]\!]_\beta(c_{\mathbf{s}}) \mid p \in \mathbf{P_G}[\mathbf{s}, \mathbf{n}]\,\}$$

Also the fixed-point counterpart of the *MOP*-solution considers for every node of $\mathbf{G}$ a pair of DFA-informations:

The *MFP*-Solution: (Basic-block Level)

$$\forall\, c_{\mathbf{s}} \in \mathcal{C}\ \forall\, \mathbf{n} \in \mathbf{N}.\ \ MFP_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathbf{n}) =_{df} (\, N\text{-}MFP_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathbf{n}),\ X\text{-}MFP_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathbf{n})\,)$$

with

$$N\text{-}MFP_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathbf{n}) =_{df} \mathsf{pre}^{\beta}_{c_{\mathbf{s}}}(\mathbf{n}) \qquad \text{and} \qquad X\text{-}MFP_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathbf{n}) =_{df} \mathsf{post}^{\beta}_{c_{\mathbf{s}}}(\mathbf{n})$$

where $\mathsf{pre}^{\beta}_{c_{\mathbf{s}}}$ and $\mathsf{post}^{\beta}_{c_{\mathbf{s}}}$ denote the greatest solutions of the following equation system:

$$
\begin{array}{lcll}
\mathbf{pre(n)} & = & \begin{cases} c_{\mathbf{s}} & \text{if } \mathbf{n} = \mathbf{s} \\ \bigsqcap \{ \mathbf{post(m)} \mid \mathbf{m} \in \mathit{pred}_{\mathbf{G}}(\mathbf{n}) \} & \text{otherwise} \end{cases} \\
\mathbf{post(n)} & = & [\![\, \mathbf{n} \,]\!]_{\beta}(\mathbf{pre(n)})
\end{array}
$$

Level 2 – Instruction Level: On this level the information must be pushed into the basic blocks to the individual instructions. Like on the basic-block level we have to distinguish between eNtry- and eXit-informations. In addition to $[\![\ ]\!]_\beta$ this requires a second local semantic functional $[\![\ ]\!]_\iota : N \to (\mathcal{C} \to \mathcal{C})$, which defines the effect of instructions, not of basic blocks (cf. Figure 3).

The *MOP*-Solution: (Instruction Level)

$$\forall\, c_{\mathbf{s}} \in \mathcal{C} \ \forall\, n \in N. \ \mathit{MOP}_{([\![\ ]\!]_\iota, c_{\mathbf{s}})}(n) =_{df} (\, N\text{-}\mathit{MOP}_{([\![\ ]\!]_\iota, c_{\mathbf{s}})}(n),\ X\text{-}\mathit{MOP}_{([\![\ ]\!]_\iota, c_{\mathbf{s}})}(n) \,)$$

with

$$
\begin{array}{lcl}
N\text{-}\mathit{MOP}_{([\![\ ]\!]_\iota, c_{\mathbf{s}})}(n) & =_{df} & \begin{cases} N\text{-}\mathit{MOP}_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathit{block}(n)) & \text{if } n = \mathit{start}(\mathit{block}(n)) \\ [\![\, p \,]\!]_\iota(N\text{-}\mathit{MOP}_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathit{block}(n))) & \text{otherwise } (p \text{ prefix-path} \\ & \text{from } \mathit{start}(\mathit{block}(n)) \\ & \text{up to } n) \end{cases} \\
X\text{-}\mathit{MOP}_{([\![\ ]\!]_\iota, c_{\mathbf{s}})}(n) & =_{df} & [\![\, p \,]\!]_\iota(N\text{-}\mathit{MOP}_{([\![\ ]\!]_\beta, c_{\mathbf{s}})}(\mathit{block}(n))) \quad (p \text{ prefix-path} \\
& & \qquad \text{from } \mathit{start}(\mathit{block}(n)) \text{ up to and including } n)
\end{array}
$$

Similarly, we get for the fixed-point counterpart.

The *MFP*-Solution: (Instruction Level)

$$\forall\, c_{\mathbf{s}} \in \mathcal{C} \ \forall\, n \in N. \ \mathit{MFP}_{([\![\ ]\!]_\iota, c_{\mathbf{s}})}(n) =_{df} (\, N\text{-}\mathit{MFP}_{([\![\ ]\!]_\iota, c_{\mathbf{s}})}(n),\ X\text{-}\mathit{MFP}_{([\![\ ]\!]_\iota, c_{\mathbf{s}})}(n) \,)$$

with

$$N\text{-}\mathit{MFP}_{([\![\ ]\!]_\iota, c_{\mathbf{s}})}(n) =_{df} \mathsf{pre}^{\iota}_{c_{\mathbf{s}}}(n) \qquad \text{and} \qquad X\text{-}\mathit{MFP}_{([\![\ ]\!]_\iota, c_{\mathbf{s}})}(n) =_{df} \mathsf{post}^{\iota}_{c_{\mathbf{s}}}(n)$$

where $\mathsf{pre}^{\iota}_{c_{\mathbf{s}}}$ and $\mathsf{post}^{\iota}_{c_{\mathbf{s}}}$ denote the greatest solutions of the following equation system:

$$
\begin{array}{lcll}
\mathbf{pre}(n) & = & \begin{cases} \mathsf{pre}^{\beta}_{c_{\mathbf{s}}}(\mathit{block}(n)) & \text{if } n = \mathit{start}(\mathit{block}(n)) \\ \mathbf{post}(m) & \text{otherwise } (m \text{ is } n\text{'s unique predecessor} \\ & \qquad \text{in } \mathit{block}(n)) \end{cases} \\
\mathbf{post}(n) & = & [\![\, n \,]\!]_\iota(\mathbf{pre}(n))
\end{array}
$$

Note that the greatest solution of the latter equation system can be computed quite efficiently by exploiting the fact that basic blocks represent straight-line code sequences. Hence, it suffices to visit the instructions of the basic blocks in their sequential order without having to visit instructions or basic blocks again. Note, however, that this requires a different implementation than for solving the first-level equation system. Moreover, inside a basic block the eNtry-information of an instruction coincides with the eXit-information of its unique predecessor. Thus, one of these informations can be dropped. This is another source of inhomogeneity between program points complicating theoretical investigations as well as implementations further.

The Coincidence and Safety Theorem can also be given for the 2-level setting of node-labeled BB-graphs. We omit this here for brevity. From the preceding presentation it should be obvious that the technical and notational details are much more complicated than for edge-labeled SI-graphs. We remark that this complexity is not restricted to theoretical investigations on correctness and precision. It directly carries over to the implementation side. One has to define and implement a DFA both on the basic-block as well as on the instruction level. This is illustrated in the following section. Note that the equation systems occurring in the following examples are specializations of the equation systems of paragraphs A) and B).

4.1.2 Availability of Terms: A Typical Application

In this section we demonstrate the impact of the choice of a specific flow-graph variant on the form of a DFA-specification considering a practically relevant analysis problem, the *availability* of terms, a representative of Hecht's famous taxonomy of DFA-problems [12]. Intuitively, a term t is available at a program point, if it has been computed on every program path reaching this point without an intervening modification.

Table 1 recalls the specification of the availability problem for *node-labeled BB-graphs*. Note that the two-level specification of the BB-approach requires a two-phase computation process. The first phase is concerned with basic blocks, the second phase with their individual instructions. For comparison, Table 2 specifies the availability problem for *node-labeled SI-graphs*. In essence, the specification reduces to the specification of the second level of its BB-counterpart.

Table 3 shows the availability specification for *edge-labeled BB-graphs*. Like its counterpart for node-labeled BB-graphs it requires a two-level specification. However, in contrast to its node-labeled counterpart, the edge-labeled version does not have to deal with entry- and exit-properties. All program points are homogeneous!

As expected, for *edge-labeled SI-graphs* the availability specification is even simpler. It is displayed in Table 4. In fact, it is the most concise and elegant one of the four variants. A fact, which applies to other DFA-problems as well. In the Appendix we demonstrate this for the problems of *very busy expressions* and *constant propagation*.

4.2 Demand for Pre- and Postprocesses or "Tricky" Formulations

After focussing on analysis in the previous section, we now concentrate on optimization. Optimizing transformations on BB-graphs demand typically for pre- and postprocesses in order to manage the technical subtleties caused by their hierarchical structure. We illustrate this by means of the busy-code-motion (BCM) transforma-

Availability for Node-labeled BB-Graphs:

Phase I: The Basic-block Level

Local Predicates: (associated with basic-block nodes)

- $\mathsf{BB\text{-}XCOMP}_\beta(t)$: β contains an instruction ι computing t, and neither ι nor any instruction of β following ι modifies an operand of t.
- $\mathsf{BB\text{-}TRANSP}_\beta(t)$: β contains no instruction modifying an operand of t.

The Equation System of Phase I:

$$\mathsf{BB\text{-}N\text{-}AVAIL}_\beta = \begin{cases} ff & \text{if } \beta = \mathbf{s} \\ \prod_{\hat\beta \in pred(\beta)} \mathsf{BB\text{-}X\text{-}AVAIL}_{\hat\beta} & \text{otherwise} \end{cases}$$

$$\mathsf{BB\text{-}X\text{-}AVAIL}_\beta = \mathsf{BB\text{-}N\text{-}AVAIL}_\beta \cdot \mathsf{BB\text{-}TRANSP}_\beta + \mathsf{BB\text{-}XCOMP}_\beta$$

Phase II: The Instruction Level

Local Predicates: (associated with single-instruction nodes)

- $\mathsf{COMP}_\iota(t)$: ι computes t.
- $\mathsf{TRANSP}_\iota(t)$: ι does not modify an operand of t.
- $\mathsf{BB\text{-}N\text{-}AVAIL}^\star$, $\mathsf{BB\text{-}X\text{-}AVAIL}^\star$: greatest solution of the equation system of Phase I.

The Equation System of Phase II:

$$\mathsf{N\text{-}AVAIL}_\iota = \begin{cases} \mathsf{BB\text{-}N\text{-}AVAIL}^\star_{block(\iota)} & \text{if } \iota = start(block(\iota)) \\ \mathsf{X\text{-}AVAIL}_{pred(\iota)} & \text{otherwise (note that } |\,pred(\iota)\,| = 1) \end{cases}$$

$$\mathsf{X\text{-}AVAIL}_\iota = \begin{cases} \mathsf{BB\text{-}X\text{-}AVAIL}^\star_{block(\iota)} & \text{if } \iota = end(block(\iota)) \\ (\mathsf{N\text{-}AVAIL}_\iota + \mathsf{COMP}_\iota) \cdot \mathsf{TRANSP}_\iota & \text{otherwise} \end{cases}$$

Tab. 1: Node-labeled BB-graphs: Availability of term t.

tion of [20] for *eliminating partially redundant expressions* in a program; the latter an optimization, which is implemented in many advanced industrial compiler systems like e.g. on the basis of [20] in the Sun SPARCompiler language systems.[4] In essence, the BCM-transformation places computations as early-as-possible in a program. This maximizes the potential of redundant code which can be eliminated by replacing the original computations of the program by references to temporaries initialized at the earliest possible program points. As proved in [20] this leads to

[4]SPARCompiler is a registered trademark of SPARC International, Inc., and is licensed exclusively to Sun Microsystems, Inc.

Availability for Node-labeled SI-Graphs:

Local Predicates: (associated with nodes)

- $\mathsf{COMP}_\iota(t)$: ι computes t.
- $\mathsf{TRANSP}_\iota(t)$: ι does not modify an operand of t.

The Equation System:

$$\mathsf{N\text{-}AVAIL}_\iota = \begin{cases} \mathit{ff} & \text{if } \iota = s \\ \prod\limits_{\hat\iota \in pred(\iota)} \mathsf{X\text{-}AVAIL}_{\hat\iota} & \text{otherwise} \end{cases}$$

$$\mathsf{X\text{-}AVAIL}_\iota = (\mathsf{N\text{-}AVAIL}_\iota + \mathsf{COMP}_\iota) \cdot \mathsf{TRANSP}_\iota$$

Tab. 2: Node-labeled SI-graphs: Availability of term t.

even *computationally optimal* results, which cannot be improved any further by means of partial redundancy elimination. In essence, the computation of the *earliest* computation points for a computation requires the computation of the set of program points, where it is *available* (i.e., where it has been computed on every program path reaching the point without an intervening modification of any of its operands), and where it is *very busy* (i.e., where it will be computed on every program continuation without a preceding modification of any of its operands). The availability analysis has been considered in detail in Section 4.1.2; the very-busyness analysis is completely dual, and is given in Appendix A. Of course, for BB-graphs the computation of availability and very busyness requires a two-level approach. In [20] this two-level approach is avoided by computing a somehow "tricky" variant of availability and very busyness, below called *up-safety* and *down-safety*, respectively. The point of the modification is that the properties computed do not hold for the "natural" entry and exit point of a BB-node (as perhaps also suggested by the predicate names in Table 5), but for a transformation-specific entry- and exit-*insertion* point inside the basic block itself, which depend on the computation pattern under consideration. This is illustrated in Figure 4. For further details see [20].[5]

Though this avoids a postprocess, it makes reasoning about the correctness of the transformation more intricate as the *MFP*-solutions computed do not coincide with the "standard" *MOP*-solutions. The Coincidence Theorem (cf. Section 4.1.1) cannot directly be applied! In addition, the BCM-transformation still relies on a preprocess eliminating partial redundancies locally inside a basic block.[6] The necessity of this local elimination is illustrated in Figure 4. The computation of $a + b$ in the fifth instruction is redundant with respect to that in the fourth instruction. It can only

[5]In essence, the entry-insertion point of a computation t is immediately in front of the first occurrence of t in the basic block, and its exit-insertion point is directly behind the last modification of one of its operands.

[6]This is not specific for BCM, but applies to every PRE-algorithm working on BB-graphs.

Availability for Edge-labeled BB-Graphs:

Phase I: The Basic-block Level

Local Predicates: (associated with basic-block edges)

- $\mathsf{BB\text{-}XCOMP}_\beta(t)$: β contains an instruction ι computing t, and neither ι nor any instruction of β following ι modifies an operand of t.
- $\mathsf{BB\text{-}TRANSP}_\beta(t)$: β contains no instruction modifying an operand of t.

The Equation System of Phase I:

$$\mathsf{BB\text{-}AVAIL}_{\mathbf{n}} = \begin{cases} \mathit{ff} & \text{if } \mathbf{n} = \mathbf{s} \\ \prod\limits_{\mathbf{m} \in pred(\mathbf{n})} \mathsf{BB\text{-}XCOMP}_{(\mathbf{m},\mathbf{n})} + \mathsf{BB\text{-}AVAIL}_{\mathbf{m}} \cdot \mathsf{BB\text{-}TRANSP}_{(\mathbf{m},\mathbf{n})} & \text{otherwise} \end{cases}$$

Phase II: The Instruction Level

Local Predicates: (associated with single-instruction edges)

- $\mathsf{COMP}_\varepsilon(t)$: instruction ι of edge ε computes t.
- $\mathsf{TRANSP}_\varepsilon(t)$: instruction ι of edge ε does not modify an operand of t.
- $\mathsf{BB\text{-}AVAIL}^\star$: greatest solution of the equation system of Phase I.

The Equation System of Phase II:

$$\mathsf{AVAIL}_n = \begin{cases} \mathsf{BB\text{-}AVAIL}^\star_{block(n)} & \text{if } n = start(block(n)) \\ (\mathsf{AVAIL}_{pred(n)} + \mathsf{COMP}_{(pred(n),n)}) \cdot \mathsf{TRANSP}_{(pred(n),n)} & \text{otherwise (note that } |\,pred(n)\,| = 1) \end{cases}$$

Tab. 3: Edge-labeled BB-graphs: Availability of term t.

be eliminated by a local inspection of the basic block.

Subsequently, Table 6 shows how the complexity of the specification of the *BCM*-transformation reduces for a setting with *node-labeled SI-graphs*. Note that down-safety and up-safety now coincide with very busyness and availability. There is no longer a need for "tricky" formulations or for any pre- or postprocesses. The counterpart of the *BCM*-transformation for *edge-labeled SI-graphs* would even be simpler due to the homogeneity of program points. We omit its presentation here for brevity. We remark, however, that edge-label modelling additionally profits from the fact that the problem of *critical edges* (see [20] for details), i.e., edges leading from nodes with more than one outgoing edge to nodes with more than one incoming edge, does not arise here.

Availability for Edge-labeled SI-Graphs:

Local Predicates: (associated with single-instruction edges)

- $\mathsf{COMP}_\varepsilon(t)$: instruction ι of edge ε computes t.
- $\mathsf{TRANSP}_\varepsilon(t)$: instruction ι of edge ε does not modify an operand of t.

The Equation System:

$$\mathsf{AVAIL}_n = \begin{cases} \mathit{ff} & \text{if } n = s \\ \prod\limits_{m \in pred(n)} (\mathsf{AVAIL}_m + \mathsf{COMP}_{(m,n)}) \cdot \mathsf{TRANSP}_{(m,n)} & \text{otherwise} \end{cases}$$

Tab. 4: Edge-labeled SI-graphs: Availability of term t.

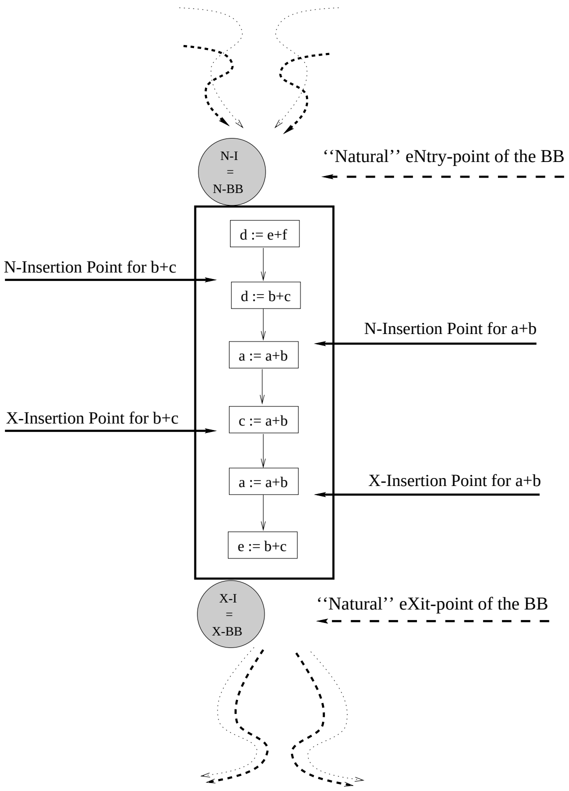

Fig. 4: eNtry-/eXit-points vs. eNtry-/eXit-insertion points of node-lab. BB-graphs.

4.3 Limited Generality

The *faint variable analysis* (cf. [11, 13, 21]) is a striking example of a practically relevant problem where it is not at all obvious of how to express it on the basic-block level. Intuitively, a variable is *faint* if there is no program continuation on which

Busy-Code-Motion for Node-labeled BB-Graphs:

1. The Up-Safety and Down-Safety Analyses

Local Predicates:

- $\mathsf{BB\text{-}NCOMP}_\beta(t)$: β contains an instruction ι computing t, which is not preceded by an instruction modifying an operand of t.
- $\mathsf{BB\text{-}XCOMP}_\beta(t)$: β contains an instruction ι computing t, and neither ι nor any instruction of β following ι modifies an operand of t.
- $\mathsf{BB\text{-}TRANSP}_\beta(t)$: β contains no instruction modifying an operand of t.

The Up-Safety Equation System:

$$\mathsf{BB\text{-}N\text{-}USAFE}_\beta = \begin{cases} \mathit{ff} & \text{if } \beta = \mathbf{s} \\ \prod_{\hat\beta \in pred(\beta)} (\mathsf{BB\text{-}XCOMP}_{\hat\beta} + \mathsf{BB\text{-}X\text{-}USAFE}_{\hat\beta}) & \text{otherwise} \end{cases}$$

$$\mathsf{BB\text{-}X\text{-}USAFE}_\beta = (\mathsf{BB\text{-}N\text{-}USAFE}_\beta + \mathsf{BB\text{-}NCOMP}_\beta) \cdot \mathsf{BB\text{-}TRANSP}_\beta$$

The Down-Safety Equation System:

$$\mathsf{BB\text{-}N\text{-}DSAFE}_\beta = \mathsf{BB\text{-}NCOMP}_\beta + \mathsf{BB\text{-}X\text{-}DSAFE}_\beta \cdot \mathsf{BB\text{-}TRANSP}_\beta$$

$$\mathsf{BB\text{-}X\text{-}DSAFE}_\beta = \mathsf{BB\text{-}XCOMP}_\beta + \begin{cases} \mathit{ff} & \text{if } \beta = \mathbf{e} \\ \prod_{\hat\beta \in succ(\beta)} \mathsf{BB\text{-}N\text{-}DSAFE}_{\hat\beta} & \text{otherwise} \end{cases}$$

2. The Transformation: Insertion and Replacement Points

Local Predicates:

- $\mathsf{BB\text{-}N\text{-}USAFE}^\star$, $\mathsf{BB\text{-}X\text{-}USAFE}^\star$, $\mathsf{BB\text{-}N\text{-}DSAFE}^\star$, $\mathsf{BB\text{-}X\text{-}DSAFE}^\star$: greatest solutions of the up-safety and down-safety equation systems of step 1.

$$\begin{aligned} \mathsf{N\text{-}INSERT}^{\mathsf{BCM}}_\beta &=_{df} \mathsf{BB\text{-}N\text{-}DSAFE}^\star_\beta \cdot \prod_{\hat\beta \in pred(\beta)} (\overline{\mathsf{BB\text{-}X\text{-}USAFE}^\star_{\hat\beta} + \mathsf{BB\text{-}X\text{-}DSAFE}^\star_{\hat\beta}}) \\ \mathsf{X\text{-}INSERT}^{\mathsf{BCM}}_\beta &=_{df} \mathsf{BB\text{-}X\text{-}DSAFE}^\star_\beta \cdot \overline{\mathsf{BB\text{-}TRANSP}_\beta} \end{aligned}$$

$$\begin{aligned} \mathsf{N\text{-}REPLACE}^{\mathsf{BCM}}_\beta &=_{df} \mathsf{BB\text{-}NCOMP}_\beta \\ \mathsf{X\text{-}REPLACE}^{\mathsf{BCM}}_\beta &=_{df} \mathsf{BB\text{-}XCOMP}_\beta \end{aligned}$$

Tab. 5: Node-labeled BB-graphs: Busy-code-motion for computation t.

Busy-Code-Motion for Node-labeled SI-Graphs:

1. The Up-Safety and Down-Safety Analyses

Local Predicates:

- $\mathsf{COMP}_\iota(t)$: ι computes t.
- $\mathsf{TRANSP}_\iota(t)$: ι does not modify an operand of t.

The Up-Safety Equation System:

$$\mathsf{N\text{-}USAFE}_\iota = \begin{cases} \mathit{ff} & \text{if } \iota = s \\ \prod_{\hat\iota \in pred(\iota)} \mathsf{X\text{-}USAFE}_{\hat\iota} & \text{otherwise} \end{cases}$$

$$\mathsf{X\text{-}USAFE}_\iota = (\mathsf{N\text{-}USAFE}_\iota + \mathsf{COMP}_\iota) \cdot \mathsf{TRANSP}_\iota$$

The Down-Safety Equation System:

$$\mathsf{N\text{-}DSAFE}_\iota = \mathsf{COMP}_\iota + \mathsf{X\text{-}DSAFE}_\iota \cdot \mathsf{TRANSP}_\iota$$

$$\mathsf{X\text{-}DSAFE}_\iota = \begin{cases} \mathit{ff} & \text{if } \iota = e \\ \prod_{\hat\iota \in succ(\iota)} \mathsf{N\text{-}DSAFE}_{\hat\iota} & \text{otherwise} \end{cases}$$

2. The Transformation: Insertion and Replacement Points

Local Predicates:

- $\mathsf{N\text{-}USAFE}^\star$, $\mathsf{X\text{-}USAFE}^\star$, $\mathsf{N\text{-}DSAFE}^\star$, $\mathsf{X\text{-}DSAFE}^\star$: greatest solutions of the down-safety and up-safety equation systems of step 1.

$$\mathsf{N\text{-}INSERT}^{\mathsf{BCM}}_\iota =_{df} \mathsf{N\text{-}DSAFE}^\star_\iota \cdot \prod_{\hat\iota \in pred(\iota)} (\overline{\mathsf{X\text{-}USAFE}^\star_{\hat\iota} + \mathsf{X\text{-}DSAFE}^\star_{\hat\iota}})$$

$$\mathsf{X\text{-}INSERT}^{\mathsf{BCM}}_\iota =_{df} \mathsf{X\text{-}DSAFE}^\star_\iota \cdot \overline{\mathsf{TRANSP}_\iota}$$

$$\mathsf{REPLACE}^{\mathsf{BCM}}_\iota =_{df} \mathsf{COMP}_\iota$$

Tab. 6: Node-labeled SI-graphs: Busy-code-motion for computation t.

it is used without a preceding modification, or if the left-hand side variable of the instruction it is used in, is faint as well. A simple example of a faint, though not dead variable, is the left-hand side occurrence of x in the statement $x := x + 1$ located inside a loop without any other occurrence of x elsewhere in the program. Below we present the specification of the faint variable analysis for both node-labeled and edge-labeled SI-graphs. We conjecture that it is impossible to express this property adequately on the BB-level. The point here is that the basic-block properties of this problem are not "really" local, but depend on the globally computed information. Hence, a basic-block analysis must be interleaved with steps for updating basic-block informations. Conceptually, this is even more complicated than the pre- and postprocesses or the "tricky" formulation of the BCM-transformation of the previous section, and destroys the two-level approach of working on BB-graphs, i.e., iterating over the BB-structure first, and inspecting them locally second.

Besides faint variable analysis, there are many other practically relevant DFA-problems like *constant propagation* (see Appendix B) or the computation of *semantically equivalent program terms* (cf. [29]), which can quite naturally and easily be expressed on the instruction level, but not on the basic-block level. In Appendix B this is illustrated for constant propagation. Note that the BB-variant given implicitly mimics the SI-variant as the effect of basic blocks is modelled by the effect of the sequential composition of its elementary instructions.

Faintness for Node-labeled SI-Graphs:

Local Predicates: (associated with single-instruction nodes)

- $\mathsf{USED}_\iota(v)$: ι uses v.
- $\mathsf{MOD}_\iota(v)$: ι modifies v.
- $\mathsf{RELV\text{-}USED}_\iota(v)$: v is a variable occurring in instruction ι forcing v to be alive (e.g., an output operation).
- $\mathsf{ASS\text{-}USED}_\iota(v)$: v is a right-hand side variable of the assignment instruction ι.

The Equation System:

$$\begin{aligned}
\mathsf{N\text{-}FAINT}_\iota(v) &= \overline{\mathsf{RELV\text{-}USED}_\iota(v)} \cdot (\mathsf{X\text{-}FAINT}_\iota(v) + \mathsf{MOD}_\iota(v)) \cdot \\
&\qquad (\mathsf{X\text{-}FAINT}_\iota(lhs_\iota) + \overline{\mathsf{ASS\text{-}USED}_\iota(v)}) \\
\mathsf{X\text{-}FAINT}_\iota(v) &= \prod_{\hat{\iota} \in succ(\iota)} \mathsf{N\text{-}FAINT}_{\hat{\iota}}(v)
\end{aligned}$$

Tab. 7: Node-labeled SI-graphs: Faintness of variable v.

---

5 Practice: Empirical Evaluation

In this section we complement our conceptual investigation by empirical results. We will see that BB-graphs do by no means compensate performance-wise for their (artificial) conceptual complexity. We compared the runtimes for different DFA-problems for *edge-labeled* BB- and SI-graphs, for programs of different size, and

Faintness for Edge-labeled SI-Graphs:

Local Predicates: (associated with single-instruction edges)

- $\mathsf{USED}_\varepsilon(v)$: instruction ι of edge ε uses v.
- $\mathsf{MOD}_\varepsilon(v)$: instruction ι of edge ε modifies v.
- $\mathsf{RELV\text{-}USED}_\varepsilon(v)$: v is a variable occurring in instruction ι of edge ε forcing v to be alive (e.g., an output operation).
- $\mathsf{ASS\text{-}USED}_\varepsilon(v)$: v is a right-hand side variable of the assignment instruction ι of edge ε.

The Equation System:

$$\mathsf{FAINT}_n(v) \;=\; \prod_{m \in succ(n)} \overline{\mathsf{RELV\text{-}USED}_{(n,m)}(v)} \cdot (\mathsf{FAINT}_m(v) + \mathsf{MOD}_{(n,m)}(v)) \cdot (\mathsf{FAINT}_m(lhs_{(n,m)}) + \overline{\mathsf{ASS\text{-}USED}_{(n,m)}(v)})$$

Tab. 8: Edge-labeled SI-graphs: Faintness of variable v.

varying average lengths of basic blocks. As expected, it turned out that (1) the average length of basic blocks and (2) the maximal chain length of the lattice of data-flow information are the key parameters for this comparison.

Figures 5, 6, and 7 show a representative profile of these results. For the problem of computing *very busy expressions* (cf. Appendix A), Figure 5 shows that there is no pay-off for BB-graphs if the average length of basic blocks is below 10 instructions, a number which is hardly exceeded in practice. Figure 6 shows an application where SI-graphs perform even better. Finally, Figure 7 illustrates the results of computing *available expressions* for a scenario, where the number of computation occurrences is small. In this practically frequent situation, the overhead for the basic-block handling is dramatically dominating.

The worst-case scenario for SI-graphs requires both large basic blocks, and large maximal chain lengths of the data-flow lattice, in order to force long iteration sequences. Both of these characteristics hardly arise in practice. E.g., Morel and Renvoise report that they never observed more than 3 iterations in their experiments [24], while Dhamdhere reports a number of 5 [6, 7]. Typical DFA-problems requiring lattices with longer chains (than e.g. bitvector problems or constant propagation), like e.g. the computation of semantically equivalent terms (cf. [29]) are beyond the scope of a BB-modelling (cf. Section 4.3).

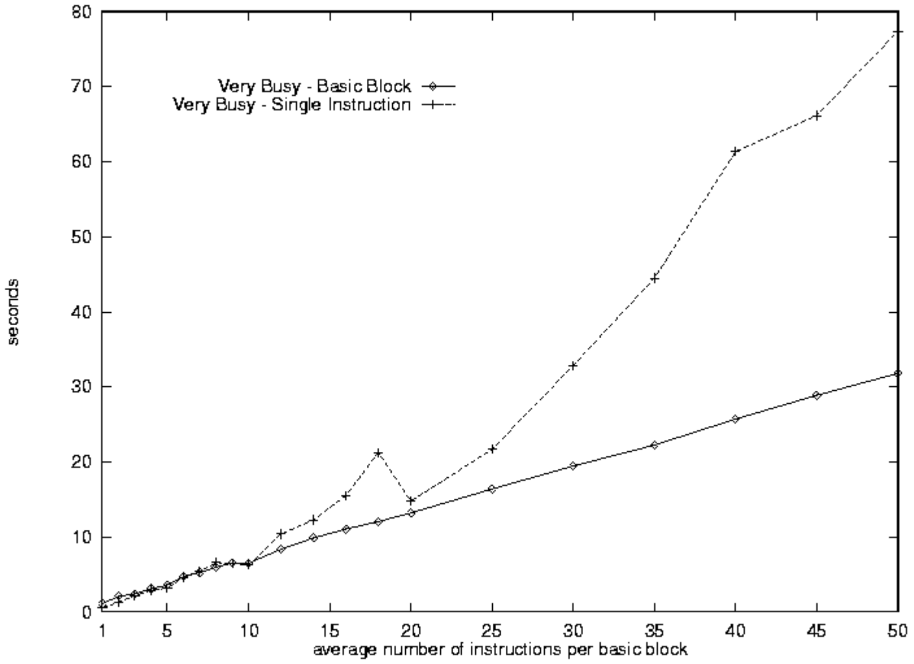

Fig. 5: No pay-off for BB-graphs with realistic BB-sizes.

6 Conclusions

For decades, BB-graphs are the state-of-the-art means for representing programs in analysis and optimization. They are considered a guarantor of high performance and broad applicability, which is believed to fairly balance the higher conceptual complexity they cause for theoretical reasoning and implementation. In this article we have systematically investigated the benefits and short-comings of the complete taxonomy of flow-graph variants. As a central result it has turned out that the severe short-comings of the currently most prominent representation is by no means compensated by its assumed benefit, namely performance. Empirical results show that the conceptually far superior SI-graphs are competetive in practice, often even superior! In fact, in everyday's life, we never experienced a situation, where the classical representation performed better. This strongly indicates that *edge-labeled SI-graphs* are the adequate representation for the considered application scenario. In fact, the experience with our DFA&OPT-generator (cf. [19]), which is based on edge-labeled SI-modeling, is extremely promising.

References

[1] J. Cai and R. Paige. "Look Ma, no hashing, and no arrays neither. In *Conference*

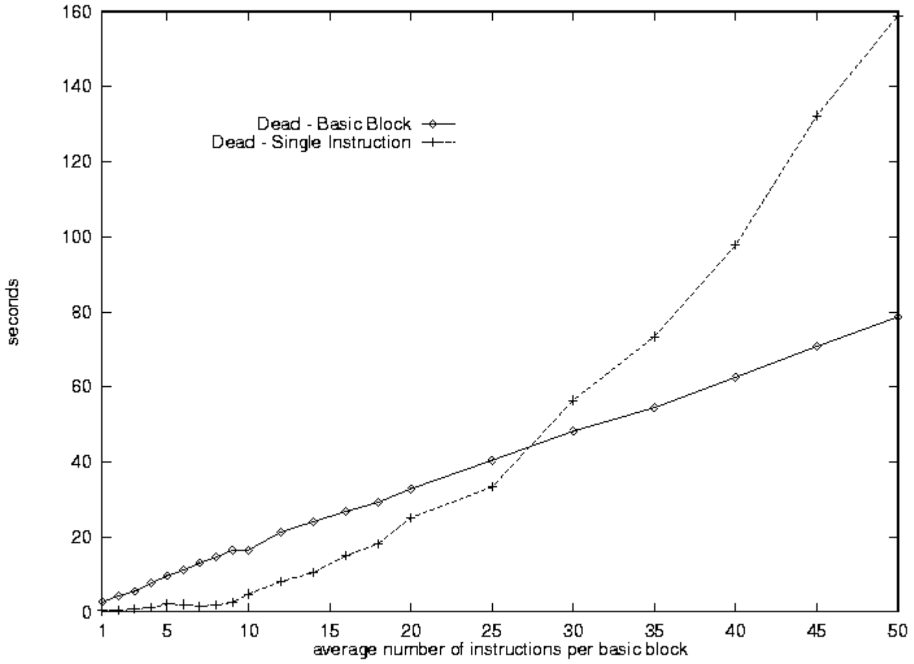

Fig. 6: A pay-off for BB-graphs at most for unrealistic large BB-sizes.

Record of the 18th Annual Symposium on Principles of Programming Languages (POPL'91) (Orlando, Florida), pages 143 – 154. ACM, New York, 1991.

[2] C. Click. Global code motion/global value numbering. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'95) (La Jolla, California)*, volume *30*,6 of *ACM SIGPLAN Notices*, pages 246–257, 1995.

[3] C. Click and K. D. Cooper. Combining analyses, combining optimizations. *ACM Transactions on Programming Languages and Systems*, 17(2):181 – 196, 1995.

[4] J. Cocke and K. Kennedy. An algorithm for reduction of operator strength. *Communications of the ACM*, 20(11):850 – 856, 1977.

[5] D. M. Dhamdhere. Characterization of program loops in code optimization. *Journal of Computer Languages*, 8(2):69 – 76, 1983.

[6] D. M. Dhamdhere. A fast algorithm for code movement optimization. *ACM SIGPLAN Notices*, 23(10):172 – 180, 1988.

[7] D. M. Dhamdhere. A new algorithm for composite hoisting and strength reduction optimisation (+ Corrigendum). *International Journal of Computer Mathematics*, 27:1 – 14 (+ 31 – 32), 1989.

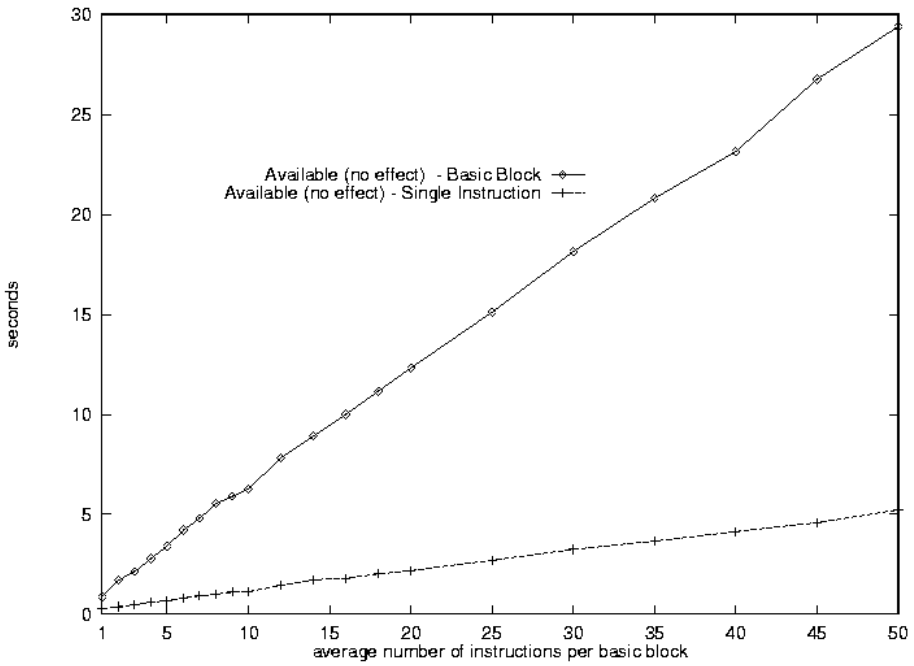

Fig. 7: The need of computing local BB-properties can be a dramatic overhead.

[8] D. M. Dhamdhere, B. K. Rosen, and F. K. Zadeck. How to analyze large programs efficiently and informatively. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation* (*PLDI'92*) (*San Francisco, California*), volume *27*,7 of *ACM SIGPLAN Notices*, pages 212 – 223, 1992.

[9] V. M. Dhaneshwar and D. M. Dhamdhere. Strength reduction of large expressions. *Journal of Programming Languages*, 3(2):95 – 120, 1995.

[10] A. Fong, J. B. Kam, and J. D. Ullman. Application of lattice algebra to loop optimization. In *Conference Record of the 2nd Annual Symposium on Principles of Programming Languages* (*POPL'75*) (*Palo Alto, California*), pages 1 – 9. ACM, New York, 1975.

[11] R. Giegerich, U. Möncke, and R. Wilhelm. Invariance of approximative semantics with respect to program transformations. In *Proceedings of the 3rd Conference of the European Co-operation in Informatics*, Informatik-Fachberichte, vol. 50, pages 1 – 10. Springer-Verlag, Heidelberg, Germany, 1981.

[12] M. S. Hecht. *Flow Analysis of Computer Programs.* Elsevier, North-Holland, 1977.

[13] S. Horwitz, A. Demers, and T. Teitelbaum. An efficient general iterative algorithm for data flow analysis. *Acta Informatica*, 24:679 – 694, 1987.

[14] S. M. Joshi and D. M. Dhamdhere. A composite hoisting-strength reduction transformation for global program optimization – Part I and Part II. *International Journal of Computer Mathematics*, 11:21 – 41, 111 – 126, 1982.

[15] J. B. Kam and J. D. Ullman. Monotone data flow analysis frameworks. *Acta Informatica*, 7:309 – 317, 1977.

[16] K. Kennedy. A survey of data flow analysis techniques. In S. S. Muchnick and N. D. Jones, editors, *Program Flow Analysis: Theory and Applications*, chapter 1, pages 5 – 54. Prentice Hall, Englewood Cliffs, New Jersey, 1981.

[17] G. A. Kildall. *Global Expression Optimization During Compilation.* PhD thesis, Univ. of Washington, Computer Science Group, Seattle, Washington, 1972.

[18] G. A. Kildall. A unified approach to global program optimization. In *Conference Record of the 1st Annual Symposium on Principles of Programming Languages* (*POPL'73*) (*Boston, Massachusetts*), pages 194 – 206. ACM, New York, 1973.

[19] M. Klein, J. Knoop, D. Koschützki, and B. Steffen. DFA&OPT-MetaFrame: A tool kit for program analysis and optimization. In *Proceedings of the 2nd International Workshop on Tools and Algorithms for the Construction and Analysis of Systems* (*TACAS'96*) (*Passau, Germany*), Lecture Notes in Computer Science, vol. 1055, pages 422 – 426. Springer-Verlag, Heidelberg, Germany, 1996.

[20] J. Knoop, O. Rüthing, and B. Steffen. Optimal code motion: Theory and practice. *ACM Transactions on Programming Languages and Systems*, 16(4):1117–1155, 1994.

[21] J. Knoop, O. Rüthing, and B. Steffen. Partial dead code elimination. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation* (*PLDI'94*) (*Orlando, Florida*), volume *29*,6 of *ACM SIGPLAN Notices*, pages 147 – 158, 1994.

[22] J. Knoop, B. Steffen, and J. Vollmer. Parallelism for free: Bitvector analyses → No state explosion! In *Proceedings of the 1st International Workshop on Tools and Algorithms for the Construction and Analysis of Systems* (*TACAS'95*) (*Aarhus, Denmark*), Lecture Notes in Computer Science, vol. 1019, pages 264 – 289. Springer-Verlag, Heidelberg, Germany, 1995.

[23] T. J. Marlowe and B. G. Ryder. Properties of data flow frameworks: A unified model. *Acta Informatica*, 27:121 – 163, 1990.

[24] E. Morel and C. Renvoise. Global optimization by suppression of partial redundancies. *Communications of the ACM*, 22(2):96 – 103, 1979.

[25] F. Mueller and D. B. Whalley. Avoiding conditional branches by code replication. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation* (*PLDI'95*) (*La Jolla, California*), volume *30*,6 of *ACM SIGPLAN Notices*, pages 56 – 66, 1995.

[26] B. K. Rosen. Data flow analysis for procedural languages. *Journal of the Association for Computing Machinery*, 26(2):322 – 344, 1979.

[27] B. K. Rosen, M. N. Wegman, and F. K. Zadeck. Global value numbers and redundant computations. In *Conference Record of the 15th Annual Symposium on Principles of Programming Languages* (*POPL'88*) (*San Diego, California*), pages 2 – 27. ACM, New York, 1988.

[28] M. Sharir and A. Pnueli. Two approaches to interprocedural data flow analysis. In S. S. Muchnick and N. D. Jones, editors, *Program Flow Analysis: Theory and Applications*, chapter 7, pages 189 – 233. Prentice Hall, Englewood Cliffs, New Jersey, 1981.

[29] B. Steffen, J. Knoop, and O. Rüthing. The value flow graph: A program representation for optimal program transformations. In *Proceedings of the 3rd European Symposium on Programming* (*ESOP'90*) (*Copenhagen, Denmark*), Lecture Notes in Computer Science, vol. 432, pages 389 – 405. Springer-Verlag, Heidelberg, Germany, 1990.

[30] B. Steffen, T. Margaria, A. Claßen, V. Braun, and M. Reitenspieß. An environment for the creation of intelligent network services. In *Annual Review of Communication, Int. Engineering Consortium Chicago* (*USA*)*, IEC,* (*invited contribution*), pages 919 – 935, 1996. Also invited contribution to the book *"Intelligent Networks: IN/AIN Technologies, Operations, Services, and Applications – A Comprehensive Report"*, IEC, 1996, pp. 287-300.

[31] B. Steffen, T. Margaria, A. Claßen, V. Braun, M. Reitenspieß, and H. Wendler. Service creation: Formal verification and abstract views. In *Proceedings of the 4th International Conference on Intelligent Networks* (*ICIN'96*) (*Bordeaux, France*), pages 96 – 101, 1996.

[32] St. W. K. Tijan and J. L. Hennessy. Sharlit — A tool for building optimizers. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation* (*PLDI'92*) (*San Francisco, California*), volume *27*,7 of *ACM SIGPLAN Notices*, pages 82 – 93, 1992.

[33] J. D. Ullman. Fast algorithms for the elimination of common subexpressions. *Acta Informatica*, 2:191 – 213, 1973.

[34] M. Wolfe. Beyond induction variables. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation* (*PLDI'92*) (*San Francisco, California*), volume *27*,7 of *ACM SIGPLAN Notices*, pages 162 – 174, 1992.

A Very Busyness

Tables 9 and 10, and Tables 11 and 12 show the specifications of the very busyness analysis for a term t for node- and edge-labeled graphs, respectively. We recall that a term t is *very busy* at a program point, if it is computed on any program continuation without a preceding modification of any of its operands (cf. [12]).

Very Busyness for Node-labeled BB-Graphs:

Phase I: The Basic-Block Level

Local Predicates: (associated with basic-block nodes)

- $\mathsf{BB\text{-}NCOMP}_\beta(t)$: β contains an instruction ι computing t, which is not preceded by an instruction modifying an operand of t.
- $\mathsf{BB\text{-}TRANSP}_\beta(t)$: β contains no instruction modifying an operand of t.

The Equation System of Phase I:

$$\mathsf{BB\text{-}N\text{-}VBE}_\beta = \mathsf{BB\text{-}NCOMP}_\beta + \mathsf{BB\text{-}X\text{-}VBE}_\beta \cdot \mathsf{BB\text{-}TRANSP}_\beta$$

$$\mathsf{BB\text{-}X\text{-}VBE}_\beta = \begin{cases} \mathit{ff} & \text{if } \beta = \mathbf{e} \\ \prod_{\hat{\beta} \in succ(\beta)} \mathsf{BB\text{-}N\text{-}VBE}_{\hat{\beta}} & \text{otherwise} \end{cases}$$

Phase II: The Instruction Level

Local Predicates: (associated with single-instruction nodes)

- $\mathsf{COMP}_\iota(t)$: ι computes t.
- $\mathsf{TRANSP}_\iota(t)$: ι does not modify an operand of t.
- $\mathsf{BB\text{-}N\text{-}VBE}^\star$, $\mathsf{BB\text{-}X\text{-}VBE}^\star$: greatest solution of the equation system of Phase I.

The Equation System of Phase II:

$$\mathsf{N\text{-}VBE}_\iota = \begin{cases} \mathsf{BB\text{-}N\text{-}VBE}^\star_{block(\iota)} & \text{if } \iota = start(block(\iota)) \\ \mathsf{COMP}_\iota + \mathsf{X\text{-}VBE}_\iota \cdot \mathsf{TRANSP}_\iota & \text{otherwise} \end{cases}$$

$$\mathsf{X\text{-}VBE}_\iota = \begin{cases} \mathsf{BB\text{-}X\text{-}VBE}^\star_{block(\iota)} & \text{if } \iota = end(block(\iota)) \\ \mathsf{N\text{-}VBE}_{succ(\iota)} & \text{otherwise (note that } |\, succ(\iota)\,| = 1) \end{cases}$$

Tab. 9: Node-labeled BB-graphs: Very busyness of term t.

Very Busyness for Node-labeled SI-Graphs:

Local Predicates: (associated with single-instruction nodes)

- $\mathsf{COMP}_\iota(t)$: ι computes t.
- $\mathsf{TRANSP}_\iota(t)$: ι does not modify an operand of t.

The Equation System:

$$\mathsf{N\text{-}VBE}_\iota \;=\; \mathsf{COMP}_\iota + \mathsf{X\text{-}VBE}_\iota \cdot \mathsf{TRANSP}_\iota$$

$$\mathsf{X\text{-}VBE}_\iota \;=\; \begin{cases} \mathit{ff} & \text{if } \iota = e \\ \prod\limits_{\hat\iota \in succ(\iota)} \mathsf{N\text{-}VBE}_{\hat\iota} & \text{otherwise} \end{cases}$$

Tab. 10: Node-labeled SI-graphs: Very busyness of term t.

B Constant Propagation: Simple Constants

We consider the problem of computing *simple constants* (cf. [18]). To this end, we consider terms $t \in \mathbf{T}$, which are inductively built from variables $v \in \mathbf{V}$, and operators $op \in \mathbf{Op}$. Their *semantics* is induced by an *interpretation* $I = (\mathbf{D}' \cup \{\perp, \top\}, I_0)$, where $\mathbf{D}'$ denotes a non-empty data domain, $\perp$ and $\top$ two new data not in $\mathbf{D}'$, and I_0 a function which associates with each 0–ary operator $c \in \mathbf{Op}$ a datum $I_0(c) \in \mathbf{D}'$ and with each n–ary operator $op \in \mathbf{Op}$, $n \geq 1$, a total function $I_0(op) : \mathbf{D}^n \to \mathbf{D}$, $\mathbf{D} =_{df} \mathbf{D}' \cup \{\perp, \top\}$, which is assumed to be strict (i.e., $I_0(op)(d_1, \ldots, d_n) = \perp$, whenever there exists a $j \in \{1, \ldots, n\}$ with $d_j = \perp$). $\Sigma = \{\, \sigma \mid \sigma : \mathbf{V} \to \mathbf{D} \,\}$ denotes the set of all *states* and $\sigma_\perp$ the distinct *start state*, which assigns $\perp$ to all variables $v \in \mathbf{V}$ (this choice of $\sigma_\perp$ reflects the fact that nothing is assumed about the context of a program being optimized). The *semantics* of terms $t \in \mathbf{T}$ is then given by the evaluation function $\mathcal{E} : \mathbf{T} \to (\Sigma \to \mathbf{D})$, which is inductively defined by:

$\forall\, t \in \mathbf{T}\ \forall\, \sigma \in \Sigma.$

$$\mathcal{E}(t)(\sigma) =_{df} \begin{cases} \sigma(x) & \text{if } t = x \in \mathbf{V} \\ I_0(c) & \text{if } t = c \text{ is a 0-ary operator} \\ I_0(op)(\mathcal{E}(t_1)(\sigma), \ldots, \mathcal{E}(t_n)(\sigma)) & \text{if } t = op(t_1, \ldots, t_n) \end{cases}$$

In the following we assume $\mathbf{D}' \subseteq \mathbf{T}$, i.e., data $d \in \mathbf{D}'$ are considered as 0-ary operators that evaluate to d.

Additionally, for every instruction $\iota \equiv (x := t)$ we define two functions

$\delta_\iota : \mathbf{T} \to \mathbf{T}$ by $\delta_\iota(s) =_{df} s[t/x]$ for all $s \in \mathbf{T}$

where $[t/x]$ stands for the simultaneous replacement of all occurrences of x by t, and $\theta_\iota : \Sigma \to \Sigma$, defined by: $\forall\, \sigma \in \Sigma\ \forall\, y \in \mathbf{V}$.

Very Busyness for Edge-labeled BB-Graphs:

Phase I: The Basic Block Level

Local Predicates: (associated with basic-block edges)

- $\mathsf{BB\text{-}NCOMP}_\beta(t)$: β contains an instruction ι computing t, and no instruction of β preceding ι modifies an operand of t.
- $\mathsf{BB\text{-}TRANSP}_\beta(t)$: β contains no instruction modifying an operand of t.

The Equation System of Phase I:

$$\mathsf{BB\text{-}VBE}_{\mathbf{n}} \;=\; \begin{cases} \mathit{ff} & \text{if } \mathbf{n} = \mathbf{e} \\ \prod\limits_{\mathbf{m} \in succ(\mathbf{n})} \mathsf{BB\text{-}NCOMP}_{(\mathbf{n},\mathbf{m})} + \mathsf{BB\text{-}VBE}_{\mathbf{m}} \cdot \mathsf{BB\text{-}TRANSP}_{(\mathbf{n},\mathbf{m})} & \text{otherwise} \end{cases}$$

Phase II: The Instruction Level

Local Predicates: (associated with single-instruction edges)

- $\mathsf{COMP}_\varepsilon(t)$: instruction ι of edge ε computes t.
- $\mathsf{TRANSP}_\varepsilon(t)$: instruction ι of edge ε does not modify an operand of t.
- $\mathsf{BB\text{-}VBE}^\star$: greatest solution of the equation system of Phase I.

The Equation System of Phase II:

$$\mathsf{VBE}_n \;=\; \begin{cases} \mathsf{BB\text{-}VBE}^\star_{block(n)} & \text{if } n = end(block(n)) \\ \mathsf{COMP}_{(n,succ(n))} + \mathsf{VBE}_{succ(n)} \cdot \mathsf{TRANSP}_{(n,succ(n))} & \text{otherwise (note that } \mid succ(n) \mid = 1) \end{cases}$$

Tab. 11: Edge-labeled BB-graphs: Very busyness of term t.

$$\theta_\iota(\sigma)(y) =_{df} \begin{cases} \mathcal{E}(t)(\sigma) & \text{if } y = x \\ \sigma(y) & \text{otherwise} \end{cases}$$

δ_ι realizes the *backward substitution* and θ_ι the *state transformation* caused by the instruction ι. Important is the following relationship between δ_ι and θ_ι, which follows by a simple inductive argument on the structure of the term $t \in \mathbf{T}$:

Lemma B.1 (Substitution Lemma)

$$\forall\, t \in \mathbf{T} \;\forall\, \sigma \in \Sigma \;\forall\, \iota \in \mathcal{I}.\; \mathcal{E}(\delta_\iota(t))(\sigma) = \mathcal{E}(t)(\theta_\iota(\sigma))$$

Very Busyness for Edge-labeled SI-Graphs:

Local Predicates: (associated with single-instruction edges)

- $\mathsf{COMP}_\varepsilon(t)$: instruction ι of edge ε computes t.
- $\mathsf{TRANSP}_\varepsilon(t)$: instruction ι of edge ε does not modify an operand of t.

The Equation System:

$$\mathsf{VBE}_n \quad = \quad \begin{cases} \mathit{ff} & \text{if } n = e \\ \prod\limits_{m \in succ(n)} \mathsf{COMP}_{(m,n)} + \mathsf{VBE}_m \cdot \mathsf{TRANSP}_{(m,n)} & \text{otherwise} \end{cases}$$

Tab. 12: Edge-labeled SI-graphs: Very busyness of term t.

Both backward substitution and state transformation can be extended to basic blocks (and, even more general, to finite paths). For each path $p = (n_1, \ldots, n_q) \in \mathbf{P}[n_1, n_q]$, we define:

- $\Delta_p : \mathbf{T} \to \mathbf{T}$ by $\Delta_p =_{df} \delta_{n_q}$, if $q = 1$, and $\Delta_{(n_1,\ldots,n_{q-1})} \circ \delta_{n_q}$, otherwise,
- $\Theta_p : \Sigma \to \Sigma$ by $\Theta_p =_{df} \theta_{n_1}$, if $q = 1$, and $\Theta_{(n_2,\ldots,n_q)} \circ \theta_{n_1}$, otherwise.

The key for proving the correctness of the constant propagation analysis defined next is the following inductive extension of the Substitution Lemma B.1.

Lemma B.2 (Generalized Substitution Lemma)

$$\forall\, t \in \mathbf{T}\ \forall\, \sigma \in \Sigma\ \forall\, \beta \in \mathcal{B}.\ \mathcal{E}(\Delta_\beta(t))(\sigma) = \mathcal{E}(t)(\Theta_\beta(\sigma))$$

We remark that the definitions and lemmas apply analogously to node- and edge-labeled graphs.

Constant Propagation for Node-lab. BB-Graphs:

Phase I: The Basic-block Level

Remark:

- $\Delta_\beta(v) =_{df} \delta_{\iota_1} \circ \ldots \circ \delta_{\iota_q}(v)$, where $\beta \equiv \iota_1; \ldots; \iota_q$.
- $\mathsf{BB\text{-}N\text{-}CP}_\beta, \mathsf{BB\text{-}X\text{-}CP}_\beta, \mathsf{N\text{-}CP}_\iota, \mathsf{X\text{-}CP}_\iota \in \Sigma$

The Equation System of Phase I:

$$\begin{aligned}
\mathsf{BB\text{-}N\text{-}CP}_\beta &= \begin{cases} \sigma_0 & \text{if } \beta = \mathbf{s} \\ \bigsqcap\{\mathsf{BB\text{-}X\text{-}CP}_{\hat\beta} \mid \hat\beta \in pred(\beta)\} & \text{otherwise} \end{cases} \\
\forall\, v \in \mathbf{V}.\ \mathsf{BB\text{-}X\text{-}CP}_\beta(v) &= \mathcal{E}(\Delta_\beta(v))(\mathsf{BB\text{-}N\text{-}CP}_\beta)
\end{aligned}$$

Phase II: The Instruction Level

Precomputed Results:

- $\mathsf{BB\text{-}N\text{-}CP}^\star$, $\mathsf{BB\text{-}X\text{-}CP}^\star$: greatest solution of the equation system of Phase I.

The Equation System of Phase II:

$$\begin{aligned}
\mathsf{N\text{-}CP}_\iota &= \begin{cases} \mathsf{BB\text{-}N\text{-}CP}^\star_{block(\iota)} & \text{if } \iota = start(block(\iota)) \\ \mathsf{X\text{-}CP}_{pred(\iota)} & \text{otherwise (note that } \mid pred(\iota) \mid = 1) \end{cases} \\
\forall\, v \in \mathbf{V}.\ \mathsf{X\text{-}CP}_\iota(v) &= \begin{cases} \mathsf{BB\text{-}X\text{-}CP}^\star_{block(\iota)}(v) & \text{if } \iota = end(block(\iota)) \\ \mathcal{E}(\delta_\iota(v))(\mathsf{N\text{-}CP}_\iota) & \text{otherwise} \end{cases}
\end{aligned}$$

Tab. 13: Node-labeled BB-graphs: Constant propagation.

Constant Propagation for Node-labeled SI-Graphs:

Remark:

- $\mathsf{N\text{-}CP}_\iota$, $\mathsf{X\text{-}CP}_\iota \in \Sigma$

The Equation System:

$$\begin{aligned}
\mathsf{N\text{-}CP}_\iota &= \begin{cases} \sigma_0 & \text{if } \iota = s \\ \bigsqcap\{\, \mathsf{X\text{-}CP}_{\hat\iota} \mid \hat\iota \in pred(\iota)\,\} & \text{otherwise} \end{cases} \\
\forall\, v \in \mathbf{V}.\ \mathsf{X\text{-}CP}_\iota(v) &= \mathcal{E}(\delta_\iota(v))(\mathsf{N\text{-}CP}_\iota)
\end{aligned}$$

Tab. 14: Node-labeled SI-graphs: Constant propagation.

Constant Propagation for Edge-lab. BB-Graphs:

Phase I: The Basic-block Level

Remark:

- $\Delta_\beta(v) =_{df} \delta_{\varepsilon_1} \circ \ldots \circ \delta_{\varepsilon_q}(v)$, where $\beta \equiv \varepsilon_1; \ldots; \varepsilon_q$.
- $\textsf{BB-CP}_{\mathbf{n}}, \textsf{CP}_n \in \Sigma$

The Equation System of Phase I:

$$\forall\, v \in \mathbf{V}.\ \textsf{BB-CP}_{\mathbf{n}} \quad = \quad \begin{cases} \sigma_0(v) & \text{if } \beta = \mathbf{s} \\ \bigsqcap\{\, \mathcal{E}(\Delta_{(\mathbf{m},\mathbf{n})}(v))(\textsf{BB-CP}_{\mathbf{m}}) \mid \mathbf{m} \in pred(\mathbf{n})\} & \text{otherwise} \end{cases}$$

Phase II: The Instruction Level

Precomputed Results:

- $\textsf{BB-CP}^\star$: greatest solution of the equation system of Phase I.

The Equation System of Phase II:

$$\forall\, v \in \mathbf{V}.\ \textsf{CP}_n \quad = \quad \begin{cases} \textsf{BB-CP}^\star_{block(n)}(v) & \text{if } n = start(block(n)) \\ \mathcal{E}(\delta_{(pred(n),n)}(v))(\textsf{CP}_{pred(n)}) & \text{otherwise (note that} \\ & \quad \mid pred(n) \mid = 1) \end{cases}$$

Tab. 15: Edge-labeled BB-graphs: Constant propagation.

Constant Propagation for Edge-labeled SI-Graphs:

Remark:

- $\textsf{CP}_n \in \Sigma$

The Equation System:

$$\forall\, v \in \mathbf{V}.\ \textsf{CP}_n \quad = \quad \begin{cases} \sigma_0(v) & \text{if } n = s \\ \bigsqcap\{\, \mathcal{E}(\delta_{(m,n)}(v))(\textsf{CP}_m) \mid m \in pred(n)\,\} & \text{otherwise} \end{cases}$$

Tab. 16: Edge-labeled SI-graphs: Constant propagation.